

Claims

What is claimed is:

- 1 1. A system for maintaining a high availability processing environment, said system comprising:
2 a network having a plurality of clusters, each cluster of the network comprising a plurality
3 of identical servers, each cluster of the network being directly connected to at least one other
4 cluster of the network, wherein each pair of clusters directly connected to each other is
5 characterized by each server in a first cluster of the pair of clusters being directly connected to at
6 least one server in a second cluster of the pair of clusters via a communication link; and
7 a control server adapted to monitor an operational status of said communication link, said
8 operational status of the communication link being that said communication link is operational or
9 non-operational, said control server being directly linked to at least one server in each cluster via
10 a communication channel between the control server and the at least one server.
- 1 2. The system of claim 1, further comprising:
2 a global dataset that includes an identification of each communication link in the network,
3 said global dataset being accessible to the control server; and
4 a local dataset specific to each cluster of the plurality of clusters, said local dataset
5 including an identification of each communication link in the network to which the servers of
6 said each cluster is coupled for flow of data out of the cluster, said global dataset being

7 accessible to the servers of said each cluster.

1 3. The system of claim 1, wherein the control server is adapted to monitor an operational status
2 of a first communication link between a first server of the first cluster and a second server of the
3 second cluster by sending a query signal to the first server, said query signal requesting the first
4 server to send a response signal to the control server indicating the operational status of the first
5 communication link, said operational status of the first communication link being that said first
6 communication link is operational or non-operational.

1 4. The system of claim 3, wherein the first server is adapted to respond to the query signal by
2 sending a prompt signal over the first communication link to the second server, said prompt
3 signal prompting the second server to send a return signal to the first server over the first
4 communication link, said return signal or absence thereof being indicative of the operational
5 status of the first communication link.

1 5. The system of claim 3, wherein the first cluster has a load balancer adapted to distribute data
2 traffic uniformly among the servers comprised by the first cluster, and wherein upon receiving
3 the response signal from the first server such that the response signal indicates that the first
4 communication link is non-operational, the control server is adapted to notify the load balancer
5 that the first communication link is non-operational.

1 6. The system of claim 3, wherein upon being notified that the first communication link is non-
2 operational, the load balancer is adapted to fail over the first server with respect to the first
3 communication link.

1 7. The system of claim 3, wherein upon receiving the response signal from the first server such
2 that the response signal indicates that the first communication link is non-operational, the control
3 server is adapted to inform a service node in the first cluster that the first communication link is
4 non-operational.

1 8. The system of claim 7, wherein upon being informed that the first communication link is non-
2 operational, the service node is adapted to make a determination of a cause of the first
3 communication link being non-operational.

1 9. The system of claim 8, wherein upon making said determination of said cause, the service
2 node is adapted to facilitate making the first communication link operational.

1 10. The system of claim 3, wherein the first cluster has a load balancer adapted to distribute data
2 traffic uniformly among the servers comprised by the first cluster, and wherein upon not
3 receiving the response signal from the first server within a predetermined period of time after
4 sending the query signal to the first server, the control server is adapted make a determination
5 that the first server is non-operational and to notify the load balancer that the first server is non-

6 operational.

1 11. The system of claim 10, wherein upon being notified that the first server is non-operational,
2 the load balancer is adapted to fail over the first server.

1 12. The system of claim 10, wherein upon not receiving the response signal from the first server
2 within a predetermined period of time after sending the query signal to the first server, the
3 control server is adapted to inform a service node in the first cluster that the first server is non-
4 operational.

1 13. The system of claim 12, wherein upon being informed that the first server is non-operational,
2 the service node is adapted to make a determination of a cause of the first server being non-
3 operational.

1 14. The system of claim 13, wherein upon making said determination of said cause the service
2 node is adapted to facilitate making the first server operational.

1 15. The system of claim 1, wherein each cluster of the plurality of clusters has a load balancer
2 adapted to distribute data traffic uniformly among the servers comprised by each cluster.

1 16. The system of claim 1, wherein the control server is adapted to receive a message from a first

2 server of the first cluster or from a load balancer of the first cluster, said message indicating that
3 an entity is non-operational, said entity being selected from the group consisting of a server of the
4 first cluster and a communication link between the first server of the first cluster and a second
5 server of the second cluster.

1 17. The system of claim 16, wherein upon receiving said message the control server is adapted
2 to inform a service node of the first cluster that the entity is non-operational.

1 18. The system of claim 1, wherein the control server is directly linked to a first server of the first
2 cluster and is not directly linked to a second server of the first cluster, wherein the first server is
3 directly connected to the second server, and wherein the control server is adapted to monitor an
4 operational status of the second server via direct communication with the first server coupled
5 with direct communication between the first server and the second server, said operational status
6 of the second server being that said second server is operational or non-operational.

1 19. The system of claim 1, wherein at least one cluster of the plurality of clusters does not have a
2 load balancer adapted to distribute data traffic uniformly among the servers comprised by the
3 first cluster.

1 20. The system of claim 1, wherein the plurality of clusters includes a web cluster of web servers,
2 an application cluster of application servers, and a database cluster of database servers, the web

3 cluster being directly connected to the application cluster, the application cluster being directly
4 connected to the database cluster, the web cluster adapted to communicate with the database
5 cluster by way of the application cluster functioning as an intermediary cluster between the web
6 cluster and the database cluster.

1 21. The system of claim 20, wherein the web cluster has a load balancer adapted to distribute
2 data traffic uniformly among the web servers comprised by the web cluster, wherein the
3 application cluster has a load balancer adapted to distribute data traffic uniformly among the
4 application servers comprised by the application cluster, and wherein the database cluster has a
5 load balancer adapted to distribute data traffic uniformly among the database servers comprised
6 by the database cluster.

1 22. The system of claim 20, wherein the web cluster has a load balancer adapted to distribute
2 data traffic uniformly among the web servers comprised by the web cluster, wherein the
3 application cluster has a load balancer adapted to distribute data traffic uniformly among the
4 application servers comprised by the application cluster, and wherein the database cluster does
5 not have a load balancer adapted to distribute data traffic uniformly among the database servers
6 comprised by the database cluster.

1 23. The system of claim 1, wherein the plurality of clusters includes a web cluster of web servers
2 and a database cluster of database servers, the web cluster being directly connected to the

3 database cluster, the web cluster adapted to directly communicate with the database cluster.

1 24. The system of claim 23, wherein the web cluster has a load balancer adapted to distribute
2 data traffic uniformly among the web servers comprised by the web cluster, and wherein the
3 database cluster has a load balancer adapted to distribute data traffic uniformly among the
4 database servers comprised by the database cluster.

1 25. The system of claim 23, wherein the web cluster has a load balancer adapted to distribute
2 data traffic uniformly among the web servers comprised by the web cluster, and wherein the
3 database cluster does not have a load balancer adapted to distribute data traffic uniformly among
4 the database servers comprised by the database cluster.

1 26. A method for maintaining a high availability processing environment, said method
2 comprising:
3 providing a network having a plurality of clusters, each cluster of the network comprising
4 a plurality of identical servers, each cluster of the network being directly connected to at least one
5 other cluster of the network, wherein each pair of clusters directly connected to each other is
6 characterized by each server in a first cluster of the pair of clusters being directly connected to at
7 least one server in a second cluster of the pair of clusters via a communication link; and
8 providing a control server adapted to monitor an operational status of said communication
9 link, said operational status of the communication link being that said communication link is
10 operational or non-operational, said control server being directly linked to at least one server in
11 each cluster via a communication channel between the control server and the at least one server.

1 27. The method of claim 26, said method further comprising:
2 providing a global dataset that includes an identification of each communication link in
3 the network, said global dataset being accessible to the control server; and
4 providing a local dataset specific to each cluster of the plurality of clusters, said local
5 dataset including an identification of each communication link in the network to which the
6 servers of said each cluster is coupled for flow of data out of the cluster, said global dataset being
7 accessible to the servers of said each cluster.

1 28. The method of claim 26, said method further comprising monitoring an operational status of

2 a first communication link between a first server of the first cluster and a second server of the
3 second cluster, said monitoring being performed by the control server, said monitoring including
4 sending a query signal to the first server, said query signal requesting the first server to send a
5 response signal to the control server indicating the status of the first communication link, said
6 operational status of the first communication link being that said first communication link is
7 operational or non-operational.

1 29. The method of claim 28, wherein the first server is adapted to respond to the query signal by
2 sending a prompt signal over the first communication link to the second server, said prompt
3 signal prompting the second server to send a return signal to the first server over the first
4 communication link, said return signal or absence thereof being indicative of the operational
5 status of the first communication link.

1 30. The method of claim 28, wherein the first cluster has a load balancer adapted to distribute
2 data traffic uniformly among the servers comprised by the first cluster, and wherein upon the
3 control server receiving the response signal from the first server such that the response signal
4 indicates that the first communication link is non-operational the method further comprises:
5 notifying the load balancer that the first communication link is non-operational, said notifying
6 being performed by the control server.

1 31. The method of claim 28, wherein upon the load balancer being notified that the first

2 communication link is non-operational the method further comprises: failing over the first server
3 with respect to the first communication link, said failing over being performed by the load
4 balancer.

1 32. The method of claim 28, wherein upon the control server receiving the response signal from
2 the first server such that the response signal indicates that the first communication link is non-
3 operational the method further comprises: informing a service node in the first cluster that the
4 first communication link is non-operational, said informing being performed by the control
5 server.

1 33. The method of claim 32, wherein upon the service node being informed that the first
2 communication link is non-operational the method further comprises: making a determination of
3 a cause of the first communication link being non-operational, said making a determination being
4 performed by the service node.

1 34. The method of claim 33, wherein upon the service node making said determination of said
2 cause the method further comprises: facilitating making the first communication link operational,
3 said facilitating being performed by the service node.

1 35. The method of claim 28, wherein the first cluster has a load balancer adapted to distribute
2 data traffic uniformly among the servers comprised by the first cluster, and wherein upon the

3 control server not receiving the response signal from the first server within a predetermined
4 period of time after sending the query signal to the first server the method further comprises:
5 making a determination that the first server is non-operational and notifying the load balancer
6 that the first server is non-operational, said making said determination that the first server is non-
7 operational and said notifying being performed by the control server.

1 36. The method of claim 35, wherein upon load balancer being notified that the first server is
2 non-operational the method further comprises: failing over the first server, said failing over being
3 performed by the load balancer.

1 37. The method of claim 35, wherein upon the control server not receiving the response signal
2 from the first server within a predetermined period of time after sending the query signal to the
3 first server the method further comprises: informing a service node in the first cluster that the
4 first server is non-operational, said informing being performed by the control server.

1 38. The method of claim 37, wherein upon the service node being informed that the first server is
2 non-operational the method further comprises: making a determination of a cause of the first
3 server being non-operational, said making said determination of said cause being performed by
4 the service node.

1 39. The method of claim 38, wherein upon the service node making said determination of said

2 cause the method further comprises: facilitating making the first server operational, said
3 facilitating being performed by the service node.

1 40. The method of claim 26, wherein each cluster of the plurality of clusters has a load balancer,
2 and wherein the method further comprises distributing data traffic uniformly among the servers
3 comprised by each cluster, said distributing being performed by the load balancer.

1 41. The method of claim 26, wherein the control server is adapted to receive a message from a
2 first server of the first cluster or from a load balancer of the first cluster, said message indicating
3 that an entity is non-operational, said entity being selected from the group consisting of a server
4 of the first cluster and a communication link between the first server of the first cluster and a
5 second server of the second cluster.

1 42. The method of claim 41, wherein upon the control server receiving the message the method
2 further comprises: informing a service node of the first cluster that the entity is non-operational,
3 said informing being performed by the control server.

1 43. The method of claim 26, wherein the control server is directly linked to a first server of the
2 first cluster and is not directly linked to a second server of the first cluster, wherein the first
3 server is directly connected to the second server, said method further comprising monitoring
4 the operational status of the second server via direct communication with the first server coupled

5 with direct communication between the first server and the second server, said monitoring being
6 performed by the control server, said operational status of the second server being that said
7 second server is operational or non-operational.

1 44. The method of claim 26, wherein at least one cluster of the plurality of clusters does not have
2 a load balancer adapted to distribute data traffic uniformly among the servers comprised by the
3 first cluster.

1 45. The method of claim 26, wherein the plurality of clusters includes a web cluster of web
2 servers, an application cluster of application servers, and a database cluster of database servers,
3 the web cluster being directly connected to the application cluster, the application cluster being
4 directly connected to the database cluster, the web cluster adapted to communicate with the
5 database cluster by way of the application cluster functioning as an intermediary cluster between
6 the web cluster and the database cluster.

1 46. The method of claim 45, wherein the web cluster has a load balancer adapted to distribute
2 data traffic uniformly among the web servers comprised by the web cluster, wherein the
3 application cluster has a load balancer adapted to distribute data traffic uniformly among the
4 application servers comprised by the application cluster, and wherein the database cluster has a
5 load balancer adapted to distribute data traffic uniformly among the database servers comprised
6 by the database cluster.

1 47. The method of claim 45, wherein the web cluster has a load balancer adapted to distribute
2 data traffic uniformly among the web servers comprised by the web cluster, wherein the
3 application cluster has a load balancer adapted to distribute data traffic uniformly among the
4 application servers comprised by the application cluster, and wherein the database cluster does
5 not have a load balancer adapted to distribute data traffic uniformly among the database servers
6 comprised by the database cluster.

1 48. The method of claim 26, wherein the plurality of clusters includes a web cluster of web
2 servers and a database cluster of database servers, the web cluster being directly connected to the
3 database cluster, the web cluster adapted to directly communicate with the database cluster.

1 49. The method of claim 48, wherein the web cluster has a load balancer adapted to distribute
2 data traffic uniformly among the web servers comprised by the web cluster, and wherein the
3 database cluster has a load balancer adapted to distribute data traffic uniformly among the
4 database servers comprised by the database cluster.

1 50. The method of claim 48, wherein the web cluster has a load balancer adapted to distribute
2 data traffic uniformly among the web servers comprised by the web cluster, and wherein the
3 database cluster does not have a load balancer adapted to distribute data traffic uniformly among
4 the database servers comprised by the database cluster.